

A method to estimate the significance of coincident gravitational-wave observations from compact binary coalescence

Kipp Cannon,^{1,*} Chad Hanna,^{2,†} and Drew Keppel^{3,4,‡}

¹*Canadian Institute for Theoretical Astrophysics, 60 St. George Street, University of Toronto, Toronto, ON M5S 3H8, Canada*

²*Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada*

³*Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-30167 Hannover, Germany*

⁴*Leibniz Universität Hannover, D-30167 Hannover, Germany*

Coalescing compact binary systems consisting of neutron stars and/or black holes should be detectable with upcoming advanced gravitational-wave detectors such as LIGO, Virgo, GEO and KAGRA. Gravitational-wave experiments to date have been riddled with non-Gaussian, non-stationary noise that makes it challenging to ascertain the significance of an event. A popular method to estimate significance is to time shift the events collected between detectors in order to establish a false coincidence rate. Here we propose a method for estimating the false alarm probability of events using variables commonly available to search candidates that does not rely on explicitly time shifting the events while still capturing the non-Gaussianity of the data. We present a method for establishing a statistical detection of events in the case where several silver-plated ($3-5\sigma$) events exist but not necessarily any gold-plated ($> 5\sigma$) events. We use LIGO data and a simulated, realistic, blind signal population to test our method.

I. INTRODUCTION

Detecting the gravitational-waves (GWs) from coalescing neutron stars and or black holes should be possible with advanced GW detectors such as LIGO, Virgo, GEO and KAGRA [1]. If the performance of past detectors is any indicator of the performance of future GW detectors, they are likely to be affected by non-Gaussian noise [2]. Coincident observations are crucial in validating the detection of GWs but it is necessary to establish the probability that the coincident event could arise from noise alone.

If the detectors' data were Gaussian and stationary, it would be straightforward to compute the false alarm probability (FAP) of a coincident event based solely on its signal-to-noise ratio (SNR) and the number of independent trials. With non-stationary, non-Gaussian data the SNR is not sufficient to describe the significance of an event and, furthermore, the distribution of detector noise is not known a priori.

Estimating false-coincident backgrounds from time delay coincidence associated with searches for GWs was first proposed for targeted compact binary coalescence GW searches in [3]. This method has been the commonest used in subsequent searches [4–14]. We present a method to estimate the false alarm probability of a GW event from coalescing compact objects without time shifts by measuring the false alarm probability distributions for non-coincident events using a set of common variables available to the searches. This greatly simplifies analysis and lends itself nicely to an online analysis environment.

This paper is organized as follows. In Sec. II we describe a formalism for ranking GW events and establishing the probability distribution for a given event's rank in noise. In Sec. III we present how to estimate the significance of a population of CBC events, which might include silver-plated (i.e. less than 5σ) events. In Sec. IV, we test our method with a mock, advanced detector search that uses four days of LIGO fifth science run (S5) data that has been recolored to have an Advanced LIGO spectrum containing a plausible, simulated, blind population of double neutron star binary mergers. We demonstrate that we can detect GWs from neutron star binaries with very low false alarm probability.

II. METHOD

GW searches for compact binary coalescence begin by matched filtering data in the detectors [15]. If peaks in SNR times series for more than one detector are consistent with the light travel time between detectors and timing errors, these peaks are considered to be a coincident event.

GW data to date have not been stationary and Gaussian [2] thus making it difficult to model the noise in GW searches. Non-stationary noise degrades the effectiveness of standard matched filter searches. For that reason additional signal consistency tests are often employed, such as explicit χ^2 tests [16, 17]. Non-stationarity occurs on several timescales. Here we are more concerned with short duration non-stationary bursts of noise called glitches for which χ^2 tests are very useful discriminators.

In this section we will present a method using common variables available to a compact binary search to estimate the FAP without relying on time shifting the detector data. Although many variables and measurements may be used, in this paper we consider two parameters: the matched filter SNR ρ_i and the χ^2 statistic χ_i^2 , which

* kipp.cannon@ligo.org

† chad.hanna@ligo.org

‡ drew.keppel@ligo.org

depend on the detector i , as well as parameters intrinsic to the source that the template describes such as mass and spin, $\bar{\theta}$. In this section, we introduce the framework for evaluating the FAP of GW candidates.

A. Ranking events

Here is our concise definition of a coincident gravitational wave search for compact binary sources. i) The search consists of D detectors. ii) We seek to find the significance of an event found in the D detectors localized in time. iii) The intrinsic parameters of the event will be unknown a priori. Our detection pipeline will measure the significance as a function of the parameters of the template waveform $\bar{\theta}$

For each detector i of a D detector network we use ρ_i and χ_i^2 to rank candidates with parameters $\bar{\theta}$ from least likely to be a gravitational wave to most likely. We use a standard likelihood ratio [18] defined as

$$\mathcal{L}(\rho_1, \chi_1^2, \dots, \rho_D, \chi_D^2, \bar{\theta}) = \frac{P(\rho_1, \chi_1^2, \dots, \rho_D, \chi_D^2, \bar{\theta} | s)}{P(\rho_1, \chi_1^2, \dots, \rho_D, \chi_D^2, \bar{\theta} | n)}, \quad (1)$$

where $P(\dots | s)$ is the probability of observing (\dots) given a signal, and $P(\dots | n)$ is the probability of observing (\dots) given noise. It is assumed that the signal distribution has been marginalized over all relevant parameters and the $\bar{\theta}$ refers only to the template waveform parameters that are measured by the pipeline. We make the simplifying assumption [19] that the likelihood can be factored into products of likelihoods from individual detectors,

$$\mathcal{L}(\rho_1, \chi_1^2, \dots, \rho_D, \chi_D^2, \bar{\theta}) \approx \prod_i^D \mathcal{L}_i(\rho_i, \chi_i^2, \bar{\theta}). \quad (2)$$

The simplification that the likelihood function can be built from these products implies statistical independence between detectors for both signals and noise. This results in a suboptimal ranking statistic. However, we can compute the FAP associated with this statistic, and in fact, it becomes much easier to do so.

B. Computing the FAP

The FAP is the probability of measuring a given \mathcal{L} if the data contains only noise. N.B., this is not the same as assessing the probability that the data contains only noise, which requires knowing the prior probabilities of both signal and noise. In constructing the FAP, $P(\mathcal{L}|n)$, we start with

$$P(\mathcal{L}, \bar{\theta} | n) = \int_{\Sigma} P(\mathcal{L}_1, \dots, \mathcal{L}_D, \bar{\theta} | n) d^{D-1}\Sigma, \quad (3)$$

where Σ is the surface of constant $\mathcal{L} = \prod_i^D \mathcal{L}_i$. From (2), we have, assuming that the likelihood values in noise are

independent between the detectors,

$$P(\mathcal{L}_1, \dots, \mathcal{L}_D, \bar{\theta} | n) = \prod_i^D P(\mathcal{L}_i, \bar{\theta} | n), \quad (4)$$

where $P(\mathcal{L}_i, \bar{\theta} | n)$ is obtained by marginalizing over ρ_i , and χ_i^2 in the single-detector terms,

$$P(\mathcal{L}_i, \bar{\theta} | n) = \int_{\sigma} P(\rho_i, \chi_i^2, \bar{\theta} | n) d\sigma, \quad (5)$$

where σ is the contour of constant \mathcal{L}_i in the $\{\rho_i, \chi_i^2\}$ surface at constant $\bar{\theta}$. Implicit in (4) and (5) is the assumption that the coincidence criteria do not depend on ρ_i, χ_i^2 or $\bar{\theta}$. Finally, $P(\mathcal{L}|n)$ is obtained by marginalizing over $\bar{\theta}$,

$$P(\mathcal{L}|n) = \int P(\mathcal{L}, \bar{\theta} | n) d\bar{\theta}. \quad (6)$$

The probability of observing an event with a likelihood value at least as large as some threshold \mathcal{L}^* is

$$P(\mathcal{L}^* | n) := P(\mathcal{L} \geq \mathcal{L}^* | n) = \int_{\mathcal{L}^*}^{\infty} P(\mathcal{L} | n) d\mathcal{L}. \quad (7)$$

A GW search will typically produce multiple coincident events during a given experiment. That means that there will be multiple opportunities to produce an event with a certain likelihood value. We are ultimately interested in the probability of getting one or more events with $\mathcal{L} \geq \mathcal{L}^*$ after all the events are considered. The probability of getting at least one such event after forming M independent coincidences [20] can be adjusted by the complement of the binomial distribution

$$\begin{aligned} P(\mathcal{L}^* | n_1, \dots, n_M) &:= 1 - \binom{M}{0} P(\mathcal{L}^* | n)^0 (1 - P(\mathcal{L}^* | n))^M \\ &= 1 - (1 - P(\mathcal{L}^* | n))^M. \end{aligned} \quad (8)$$

This is the FAP at \mathcal{L}^* in an experiment that yielded M coincident events. In what follows, we will drop the explicit n_1, \dots, n_M notation and simply use n where it is assumed that we have corrected for the number of trials.

III. GW EVENTS AS A POISSON DISTRIBUTION

Historically, GW experiments have used rates to rank events [10–14]. Assuming that the likelihood function is independent of time over the duration of the experiment (or can be approximated as such) we can cluster the most significant events of the search over a duration longer than the correlation induced by the filter and we might expect the events arising from noise to obey Poisson statistics. In what follows we assume that in fact this is the case and connect our estimation of FAP with

the false alarm rate (FAR) often quoted in gravitational wave searches for compact binary coalescence.

For a Poisson process with mean λ , the probability of observing N or more events is given by the survival function

$$P(N|\lambda) = 1 - e^{-\lambda} \sum_{i=0}^{N-1} \frac{\lambda^i}{i!}. \quad (9)$$

Using $P(\mathcal{L}^*)$ from (8), setting $N = 1$, and solving for λ tells us the mean number of noise events with $\mathcal{L} \geq \mathcal{L}^*$,

$$\lambda(\mathcal{L}^*) = -\ln[1 - P(\mathcal{L}^*)]. \quad (10)$$

The quantity, inverse false alarm rate, is given by $\text{IFAR} = T/\lambda$, where T is the observation time of the experiment.

If \mathcal{L}_N^* is the likelihood of the N^{th} most significant event, the number of background events expected with $\mathcal{L} \geq \mathcal{L}_N^*$ is

$$\lambda(\mathcal{L}_N^*) = -\ln[1 - P(\mathcal{L}_N^*)]. \quad (11)$$

Since we observed N events with $\mathcal{L} \geq \mathcal{L}_N^*$, the probability of having produced at least this many events is found by substituting (11) into (9),

$$P(N|\lambda(\mathcal{L}_N^*)) = 1 - e^{-\lambda(\mathcal{L}_N^*)} \sum_{i=0}^{N-1} \frac{\lambda(\mathcal{L}_N^*)^i}{i!}. \quad (12)$$

A population of events can collectively be more significant than the single most significant event alone. Indeed, population analyses have previously been employed in looking for GW signals associated with gamma-ray bursts (GRBs). For example, a Student-T test was proposed in [21] to test for deviations in the cross-correlation of detectors' output preceding a set of times associated with GRBs (i.e., *on-source* times) when compared to other *off-source* times not associated with GRBs, a binomial test was employed in [22, 23] using the $X\%$ most significant events to test for excess numbers of events at their associated FAPs, a Kolmogorov test was used in [24] to look for deviations from isotropy in GRB direction based on the directional sensitivity of the bar detectors, and a Mann-Whitney U (or Mann-Whitney-Wilcoxon) test was performed in [25] to test if the all the FAPs associated with the on-source events of the GRBs were on average smaller than the expected distribution given by the off-source events, as would be the case if the average significance were elevated due to the presence of GWs in the on-source events.

As noted in [22, 23], seeking significance by considering different choices of population diminishes the significance of each on account of the trials that have been conducted. We control this by restricting ourselves to considering only populations consisting of contiguous sets of events that include the most significant, and are limited to a maximum size N_{max} where N_{max} is the rank of the most significant event at whose ranking statistic (IFAR) value the expected number of background events was greater

than 1. There are N_{max} choices of population possible, so we incur that cost from the number of trials, modifying the FAP in (12) to

$$P(N|N_{\text{max}}) = 1 - (1 - P(N|\lambda(\mathcal{L}_N^*)))^{N_{\text{max}}}. \quad (13)$$

IV. EXAMPLE

We have applied these techniques to a mock search for GWs from binary neutron stars in four days of S5 LIGO data that has been recolored to match the Advanced LIGO design spectrum [26][27]. This provides a potentially realistic data set that contains glitches from the original LIGO instruments. A population of neutron star binaries was added at a rate of 4 / Mpc³ / Myr, (see [1] for the expected rates.) We self-blinded the signal parameters with a random number generator.

Our analysis targeted compact binary systems with component masses between 1.2 and 2 M_{\odot} . We used 3.5 post-Newtonian order stationary phase approximation templates to cover the parameter space with a 97% minimal match [28] by neglecting the effects of spin in the waveform models [29]. This required $\sim 15,000$ templates. We started the matched filter integrals at 15 Hz and extended the integral to the innermost stable circular orbit frequency. The analysis gathered the data, whitened it, filtered it, identified events in the single detectors, found coincidences and ranked the events by their joint likelihoods. The filtering algorithm is described in [30].

The previous section described our method for estimating the significance of events but did not describe many details of how the calculation is done in practice. We will point out a few of those details now.

The numerator of (1) is evaluated by assuming the signals follow their expected distribution in Gaussian noise. We note that this is a reasonable assumption because detections are likely to come from periods of relatively stationary and Gaussian data. Note that the expectation for ρ can be obtained by assuming that sources are distributed uniformly in space. The expectation for the χ^2 of a signal can be found in [16].

The denominator of (1) is found by explicitly histogramming the single detector events that are not found in coincidence. By excluding coincident events we lower the chance that a gravitational wave will bias the noise distribution of the likelihoods. In general the histogramming will suffer from finite statistics and “edge” effects. We generate the histograms at a finer resolution than required to track the likelihood and then apply a Gaussian smoothing kernel with a width characteristic of the uncertainty in ρ .

We are unable to collect enough statistics to fully resolve the tail of the background ρ distribution. Thus, we add a prior distribution into the background statistics that models the ρ falloff as expected from a 2 degree of freedom matched filter in Gaussian noise, i.e. $p(\rho|n) \propto \exp[-\rho^2/2]$. This helps ensure that the like-

likelihood contours increase as a function of ρ at large ρ . At some point the probability of getting a given value of ρ , χ^2 becomes smaller than double precision float epsilon. We extend the background distribution above a given value of ρ with a polynomial in ρ that falls off faster than the signal distribution (which is $\propto \rho^{-4}$) but is shallow enough to prevent numerical problems. In both cases the point of the prior is not to influence the ranking of typical events but rather to make the calculations more numerically well-behaved. The prior is added so that the total probability amounts only to a single event in each detector. Thus the background (as billions of events are collected) quickly overwhelms the prior except for at the edges where there is no data. The point where the calculation is no longer based on having at least 1 actual event in background is important since it will effectively mark the limiting FAP. More discussion of that point follows.

In Fig. 1 we show some of the intermediate data used in estimating the significance of events in our example. Namely, we show the individual likelihood contours for ρ and χ^2 described in (1) in the H1 and L1 instruments for signals with a chirp mass consistent with a neutron star binary ($1.2 M_\odot$) in Subfig. 1(a) and 1(b) respectively. The probability of getting an event with a likelihood greater than \mathcal{L}^* after M trials for the H1 and L1 instruments (13) is shown in Fig. 1(c). Our ability to measure $P(\mathcal{L}^*|n)$ is limited by the number of events that we collect in our background estimate. The shaded region shows the \sqrt{N} error region found by assuming Poisson errors on the number of events that went into computing a given point on the curve. We have indicated the FAP at which there ceases to be more than 1 event collected in the background by a dashed line. The dashed line shows the $P(\mathcal{L}^*|n)$ has background events to $\mathcal{P} := 7 \times 10^{-5}$ which is nearly the FAP required for a 4σ detection. Below the dashed line the FAP estimate is dominated by the Gaussian smoothing kernel applied to the planes in Figs. 1(a) and 1(b). We believe that it is reasonable to trust the FAP estimate beyond the single background event limit but note that 5σ level confidence can still be reached without extrapolation with tighter coincidence criteria. Tighter coincidence criteria would reduce the trials factor and permit higher significances to be estimated. The best way to do this is to demand that three or more detectors see an event. In our example a third detector would lower the trials factor by ~ 100 , which would shift the limiting FAP, \mathcal{P} to $\sim 7 \times 10^{-7}$. It is worth mentioning that the background events and number of independent trials are accumulated at the same rate. Thus one cannot decrease the limiting FAP by collecting more data.

After assigning the FAP to events we also assign a FAR according to (10). This allows us to produce the standard IFAR plot commonly produced in recent searches for compact binaries [10, 14] without having relied on time shifting the detector events to estimate the background. This is shown in Fig. 2(a).

The IFARs of the most significant events that came out




FIG. 1. Figure 1(a) and 1(b) show the likelihoods \mathcal{L}_{H1} , \mathcal{L}_{L1} as a function of ρ and χ^2 for H1 and L1 respectively for templates with masses consistent with neutron star binaries ($1.2-2 M_\odot$). \mathcal{L}_{H1} , \mathcal{L}_{L1} appear as the right-hand-side of (2). Lighter colors refer to higher likelihood values. Figure 1(c) shows the probability of having obtained a given value of likelihood \mathcal{L}^* or greater from noise as defined in (8) after M trials (where M is the number of independent coincidences formed. In this example $M = 6 \times 10^4$.)




FIG. 2. Fig. 2(a) is a standard IFAR plot where the shaded regions correspond to the “ 1σ ” through “ 7σ ” regions computed using the survival function and point percent function associated with the Poisson distribution. This is used to determine where to stop the accumulation of events for the population statement. Fig. 2(b) shows the FAP associated with each of the individual events in the population we are considering as well as the FAP of obtaining the running N loudest events without restricting the FAP to be greater than $\mathcal{P} = 7 \times 10^{-5}$. Also shown are the same traces obtained after restricting the FAPs to be greater than \mathcal{P} .

of this search in Fig. 2(a) can be identified as the long tail in the observed events distribution. The top event has a significance greater than 5σ , the level necessary for claiming the detection of GWs. The second loudest event has a significance greater than 4σ . Both events surpass the single background event limit \mathcal{P} . If restricted to this limit then both events are nearly 4σ .

Applying the population procedure we have put forth in Sec. III, we produced a more significant statement about the presence of GWs beyond that of the loudest event. This effect is mostly attributed to the similar significance of the top two events. This could happen in a real analysis in two ways 1) Nature could just provide such a set of events as in this example 2) both events exhaust our ability to measure significance and we must place an upper bound on the FAP. The latter case, al-

though somewhat artificial, could still play an important role in analysis, especially if one is unable to confidently declare a single 5σ event but finds two or more events with 3 or 4σ . With our example analysis the combination of the two loudest events was a 5σ excursion even after restricting the FAP of both events to be \mathcal{P} . After examining the signal population we found that both candidates were separately associated with signal injections.

V. CONCLUSION

We have provided a method for estimating the significance of GWs from compact binary coalescence using measurements of single instrument populations of ρ and χ^2 as a function of the template waveform intrinsic parameters. We demonstrated our method with mock Advanced LIGO data derived from initial LIGO data including a realistic population of compact binary merger signals and glitches. We found that between our two loudest events we were able to establish detection at greater than 5σ confidence. Both of the loudest two events exhausted the \mathcal{P} ($\sim 4\sigma$) background estimate, but the extrapolated FAP of the loudest event exceeded 5σ on its own. Both of the loudest events were associated with the blind signal population introduced into the data and the remaining events were consistent with the expectation from background.

ACKNOWLEDGMENTS

Research at Perimeter Institute is supported through Industry Canada and by the Province of Ontario through the Ministry of Research & Innovation. KC was supported by the National Science and Engineering Research Council, Canada. DK was supported from the Max Planck Gesellschaft. This work has LIGO document number ligo-p1200031. The authors wish to thank Rahul Biswas, Patrick Brady, Gabriela González, Jordi Burguet-Castell, Sarah Caudill and Ruslan Vaulin for many fruitful discussions. The authors also thank John Whelan for a careful review of this work. We thank the UW-Milwaukee Center for Gravitation and Cosmology and the LIGO Laboratory for use of its computing facility to make this work possible through NSF grants PHY-0923409 and PHY-0600953. We thank the LIGO Scientific Collaboration and the Virgo Scientific Collaboration for providing us with the data to test the methods described in this work.

Appendix A: Numerical Considerations

1. Equation (8)

As the duration of the experiment increases, the numerical evaluation of (8) using fixed-precision floating

point numbers becomes challenging. In this limit, the per-trial false-alarm probability of interesting events is very small and the number of trials is very large. Using double-precision floating-point numbers, when the number of trials gets larger than about 10^{10} , FAPs of 10^{-6} and 0 become indistinguishable, and as the number of coincidences that are recorded increases further “ 4σ ” and “ 5σ ” events cannot be differentiated — it is no longer possible to make detection claims. The following procedure can be used to evaluate (8) for all $P(\mathcal{L}^*|n)$ and M . If $MP(\mathcal{L}^*|n) < 1$ the Taylor expansion of (8) about $P(\mathcal{L}^*|n) = 0$ converges quickly.

$$1 - (1 - P)^M = MP - (M^2 - M)\frac{P^2}{2} + (M^3 - 3M^2 + 2M)\frac{P^3}{6} - (M^4 - 6M^3 + 11M^2 - 6M)\frac{P^4}{24} + \dots$$

$$= \sum_{i=0}^{\infty} -1^i \frac{P^{(i+1)}}{(i+1)!} [(M-0)(M-1)\dots(M-i)]. \quad (\text{A1})$$

The last form yields a recursion relation allowing subsequent terms in the series to be computed without explicit evaluation of the numerator and denominator separately (which, otherwise, would quickly overflow): if the $(i-1)$ th term is X , the i th term in the series is $X \frac{i-M}{i+1} P$.

If $MP(\mathcal{L}^*|n) \geq 1$ the Taylor series still converges (in fact, as long as the number of trials M is an integer the series is exact in a finite number of terms) but the series is numerically unstable: the terms alternate sign and one must rely on careful cancellation of large numbers to obtain an accurate result. In this regime the expression’s value is close to 1, so $(1 - P)^M$ is small. If P is small, we can write

$$1 - (1 - P)^M = 1 - e^{M \ln(1-P)} \quad (\text{A2a})$$

and then the Taylor expansion of $M \ln(1 - P)$ about $P = 0$ converges quickly,

$$M \ln(1 - P) = -MP \left(1 + \frac{P}{2} + \frac{P^2}{3} + \dots \right). \quad (\text{A2b})$$

Altogether, the algorithm for evaluating (8) is: if

$MP(\mathcal{L}^*|n) < 1$ use (A1) computed via the recursion relation; otherwise if $P(\mathcal{L}^*|n) < 0.125$ use (A2); otherwise evaluate (8) directly using normal floating point operations. The threshold of $P(\mathcal{L}^*|n) < 0.125$ for using (A2) is found empirically, the results are not sensitive to the choice of this number.

2. Equation (10)

The evaluation of (10) for events that are interesting as detection candidates after an experiment is concluded is straight-forward using double-precision floating-point arithmetic. In this regime, $P(\mathcal{L}^*|n) \sim 10^{-5}$, and there is plenty of numerical dynamic range available. However, the practical use of (10) is in its ability to identify “once a day” or “once an hour” events for the purpose of providing alerts to the transient astronomy community. After just one day, 24 “once an hour” background events are expected, and their FAP — the probability of observing at least one such event from a Poisson process you expect to have produced 24 — is 0.9999999999622486. After 37 events are expected, double-precision numbers can no longer be used to differentiate those events’ FAPs from 1; that is, (10) can only assign reliable false-alarm rates to the 30 or so most significant background events in any experiment.

This problem is addressed by not computing the expected number of events, $\lambda(\mathcal{L}^*)$, from the false-alarm probability, $P(\mathcal{L}^*)$, as shown in (10), but by first going back and rewriting (7) and (8) as

$$1 - P(\mathcal{L}^*|n_1, \dots, n_M) = \left(\int_0^{\mathcal{L}^*} P(\mathcal{L}|n) d\mathcal{L} \right)^M, \quad (\text{A3})$$

from which we can rewrite (10) as

$$\lambda(\mathcal{L}^*) = -M \ln \int_0^{\mathcal{L}^*} P(\mathcal{L}|n) d\mathcal{L}. \quad (\text{A4})$$

This form of the expression presents no challenges to its evaluation using double-precision floating point arithmetic.

-
- [1] J. Abadie *et al.* (The LIGO Scientific Collaboration), *Class. Quantum Grav.* **27**, 173001 (2010)
 - [2] J. Slutsky, L. Blackburn, D. A. Brown, L. Cadonati, J. Cain, M. Cavagli, S. Chatterji, N. Christensen, M. Coughlin, S. Desai, G. Gonzalez, T. Isogai, E. Katsavounidis, B. Rankins, T. Reed, K. Riles, P. Shawhan, J. R. Smith, N. Zotov, and J. Zweizig, *Class. Quantum Grav.* **27**, 165023 (2010)
 - [3] B. Abbott *et al.* (The LIGO Scientific Collaboration),

- Phys. Rev. D* **69**, 122001 (2004), <http://arxiv.org/abs/gr-qc/0308069>
- [4] B. Abbott *et al.* (The LIGO Scientific Collaboration), *Phys. Rev. D* **72**, 082001 (2005), <http://arxiv.org/abs/gr-qc/0505041>
- [5] B. Abbott *et al.* (The LIGO Scientific Collaboration), *Phys. Rev. D* **72**, 082002 (2005)
- [6] B. Abbott *et al.* (The LIGO Scientific Collaboration), *Phys. Rev. D* **73**, 062001 (2006)

- [7] T. L. S. Collaboration and the TAMA Collaboration (The LIGO Scientific Collaboration), Phys. Rev. D **73**, 102002 (2006)
- [8] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **78**, 042002 (2008)
- [9] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **77**, 062002 (2008), <http://arxiv.org/abs/0704.3368>
- [10] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **79**, 122001 (2009), <http://arxiv.org/abs/0901.0302>
- [11] B. P. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **80**, 047101 (Aug 2009), <http://link.aps.org/doi/10.1103/PhysRevD.80.047101>
- [12] J. Abadie *et al.* (The LIGO Scientific Collaboration and The Virgo Collaboration), Phys. Rev. D **82**, 102001 (Nov 2010), <http://link.aps.org/doi/10.1103/PhysRevD.82.102001>
- [13] J. Abadie *et al.* (The LIGO Scientific Collaboration and The Virgo Collaboration), Phys. Rev. D **83**, 122005 (Jun 2011), <http://link.aps.org/doi/10.1103/PhysRevD.83.122005>
- [14] J. Abadie *et al.* (The LIGO Scientific Collaboration and The Virgo Collaboration)(2012), <http://arxiv.org/abs/1111.7314>
- [15] B. A. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton(2005), gr-qc/0509116
- [16] B. Allen, Phys. Rev. D **71**, 062001 (2005)
- [17] C. Hanna, *Searching for gravitational waves from binary systems in non-stationary data*, Ph.D. thesis (2008)
- [18] P. Gregory, *Bayesian Logical Data Analysis for the Physical Sciences* (Cambridge University Press, 2005)
- [19] K. C. Cannon, Class. Quantum Grav. **25**, 105024 (2008)
- [20] In practice it can be difficult to know if the coincidences formed are independent, however as long as they are related to the true number of independent trials by an overall scaling, one can adjust the number so that it agrees with the observed rate of coincidences for low significance events. This works because GWs are very rare and true signals will vastly overwhelm the false positives that a pipeline produces at high FAP. Thus the bias in calibrating M to high FAP events is very small.
- [21] L. S. Finn, S. D. Mohanty, and J. D. Romano, Phys. Rev. D **60**, 121101 (Nov 1999), <http://link.aps.org/doi/10.1103/PhysRevD.60.121101>
- [22] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **77**, 062004 (Mar 2008), <http://link.aps.org/doi/10.1103/PhysRevD.77.062004>
- [23] B. P. Abbott *et al.*, The Astrophysical Journal **715**, 1438 (2010), <http://stacks.iop.org/0004-637X/715/i=2/a=1438>
- [24] P. Astone, D. Babusci, M. Bassan, P. Carelli, E. Coccia, C. Cosmelli, S. D’Antonio, V. Fafone, F. Frontera, G. Giordano, C. Guidorzi, A. Marini, Y. Minenkoy, I. Modena, G. Modestino, A. Moleti, E. Montanari, G. V. Pallottino, G. Pizzella, L. Quintieri, A. Rocchi, F. Ronga, L. Sperandio, R. Terenzi, G. Torrioli, and M. Visco, Phys. Rev. D **71**, 042001 (Feb 2005), <http://link.aps.org/doi/10.1103/PhysRevD.71.042001>
- [25] J. Abadie *et al.*, The Astrophysical Journal **715**, 1453 (2010), <http://stacks.iop.org/0004-637X/715/i=2/a=1453>
- [26] The LIGO Scientific Collaboration, “Advanced LIGO anticipated sensitivity curves,” (2009), <https://dcc.ligo.org/cgi-bin/private/DocDB/ShowDocument?docid=T0900288>
- [27] Specifically the zero-detuned, high-power noise curve was used
- [28] B. J. Owen, Phys. Rev. D **53**, 6749 (1996)
- [29] A. Buonanno, B. Iyer, E. Ochsner, Y. Pan, and B. Sathyaprakash, Phys. Rev. D **80**, :084043 (2009)
- [30] K. Cannon, R. Cariou, A. Chapman, M. Crispin-Ortuzar, N. Fotopoulos, M. Frei, C. Hanna, E. Kara, D. Keppel, L. Liao, S. Privitera, A. Searle, L. Singer, and A. Weinstein, The Astrophysical Journal **748**, 136 (2012), <http://stacks.iop.org/0004-637X/748/i=2/a=136>